



LQCD-ext II Project Change Request CR16-02

Robert D. Kennedy
Associate Contractor Project Manager

DOE FY2016 Annual Progress Review
Thomas Jefferson National Accelerator Facility
June 28-29, 2016

FY16 Change Requests

- ▶ [CR16-01 Three-Site Cluster Hosting](#): Add cluster-hosting at BNL to the baseline project plan already hosting clusters at FNAL and TJNAF
 - *Upcoming presentation*
- ▶ [CR16-02 Unified Performance Goals](#): Replace distinct project KPIs for conventional and accelerated computing with KPIs that do not distinguish by hardware architecture
 - *This presentation*
 - Background and Motivation
 - Proposed Solution
- ▶ [These CRs are independent of each other.](#)
 - While we advocate that both be approved, neither depends on the other's approval.

Background and Motivation

Excerpt from PEP Appendix D

▶ Baseline: Separate KPI's

◦ "Conventional Computing"

- Traditional x86 CPUs

◦ "Accelerated Computing"

- GPUs on Accelerator Cards

ID	Fiscal Year	Measurement Category	Measurement Indicator	Target
7	2016	Effectiveness	Additional computing resources deployed by the project, expressed as an average of the HISQ and DWF algorithm performances in TFlops. – <u>Conventional</u> Resources	≥10 TF
8	2016	Effectiveness	Additional computing resources deployed by the project, expressed as an average of the HISQ and DWF algorithm performances in TFlops. – <u>Accelerated</u> Resources	≥39 TF
9	2016	Scientific Program Support	TF-Yrs delivered towards the completion of the Scientific Program – <u>Conventional</u> Resources	68 TF-Yrs
10	2016	Scientific Program Support	TF-Yrs delivered towards the completion of the Scientific Program – <u>Accelerated</u> Resources	67 TF-Yrs

- ▶ We have built acquisition plans assuming a purchase of 50% CPU + 50% GPU by \$ (usually choose 60–40 by \$ in practice). **This has served us well for years.**
- ▶ New technologies have been emerging that do not neatly fit into either category such as MIC (eg. Intel Xeon Phi KNL). **Now, they are here, in production.**
- ▶ USQCD demands may favor one technology over another such that a 50–50 buy can no longer adjust our hardware portfolio to meet demand. **Demand is evolving.**
- ▶ USQCD demand for larger clusters (node count, memory footprint) is at odds with splitting each acquisitions 50–50 between technologies. **Algorithms are evolving.**
- ▶ Goal: Adjust hardware portfolio to maximize Science from application portfolio.
 - The CPU/GPU distinction at KPI level may constrain project to a sub-optimal h/w portfolio choice
 - Good distinction for tracking and forecasting demand and support allocation process, but for project KPI's?

Proposed Solution

▶ Proposed: Roll-up KPI's

- New Value = Conventional + Accelerated

Excerpt from proposed PEP Appendix D

ID	Fiscal Year	Measurement Category	Measurement Indicator	Target
7-8	2016	Effectiveness	Additional computing resources deployed by the project, expressed as an average of the HISQ and DWF algorithm performances in TFlops. - ALL Resources	≥49 TF
9-10	2016	Scientific Program Support	TF-Yrs delivered towards the completion of the Scientific Program - ALL Resources	135 TF-Yrs

- ▶ Use established conversion: 1 TFlop-Year = 1 effective TFlop-Year
 - In use by Scientific Program Committee for requests capable of running on either technology
- ▶ Cost, Schedule Impact: Zero. Performance Impact: >= Zero. Potentially positive.
 - Project can explore new “advanced” technologies which may potentially be more cost-effective.
 - We *continue* to track and forecast computing by technologies to meet application demands.
- ▶ Change Level per Project Change Process
 - Level 3: Federal Project Director approval required
 - “modification to the technical performance baseline defined in a Level-1 milestone”
- ▶ Risk Assessment
 - Project purchases more of X technology, less of Y technology than is optimal for Science
 - This risk exists independent of this CR. Already a process in place to mitigate this as described in PEP, Acquisition Strategy, Risk Register.
 - We adjust annual acquisitions to better match the portfolio of LQCD computing to the evolving needs of the USQCD software portfolio.

Supporting Slides

CR16-02 Affected Project Documents

▶ Affected Project Documents

- *Project Execution Plan: 1 each for CR16-01 approved or not approved cases*
 - Updated boilerplate. Includes FY15 PEP Addendum. Conventional and Accelerated Computing KPIs are combined.
- *Performance Forecast: Update to KPI calculation to support this CR.*
 - We will continue to track distinct Conventional and Accelerated Computing capabilities so long as this is useful for the SPC to use in deciding allocations. We may add new technology categories if the project and/or USQCD find them useful to support allocations.

Between Baseline PEP and Today

- ▶ [FY15 PEP Addendum](#): Split existing Site Manager role into (a) Site Manager and (b) Site Architect
 - Site Manager: Oversees Site Operations and Management at Site
 - Site Architect: Oversees Site Strategic Planning, Design & Architecture Topics
 - Why: Better capture our actual practice. Improve communications (interaction with SPC) and succession planning by distinguishing between operations and architecture responsibilities
 - Change Control Level = “0”
 - It does not trigger any thresholds. Since it does involve changes to PEP text, it was treated similar to a formal Change Request.
 - Adds no new responsibilities to Site Managers, simply divides existing responsibilities among two roles. (*See FY16 project org chart*).
 - Approved by Contract Project Manager: July 9, 2015
 - DOE Federal Project Directory informed and concurred: Sep 9, 2015
 - [This was captured in a separate “PEP Addendum” which we agreed to roll into the official PEP document with the next update to the baseline PEP.](#)
 - With either CR16 change being approved, this would be done.